

AS-CRED: Reputation Service for Trustworthy Inter-domain Routing

Krishna Venkatasubramanian

Computer and Information Science
University of Pennsylvania

ONR MURI N00014-07-1-0907
Review Meeting
June 10, 2010

Report Documentation Page			<i>Form Approved OMB No. 0704-0188</i>	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE 10 JUN 2010	2. REPORT TYPE	3. DATES COVERED 00-00-2010 to 00-00-2010		
4. TITLE AND SUBTITLE AS-CRED: Reputation Service for Trustworthy Inter-domain Routing (BRIEFING CHARTS)			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Pennsylvania, Computer and Information Science, 3451 Walnut St, Philadelphia, PA, 19104			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES MURI Review, June 2010. U.S. Government or Federal Rights License				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 26
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified		

Overview

Border Gateway Protocol

Problems with BGP

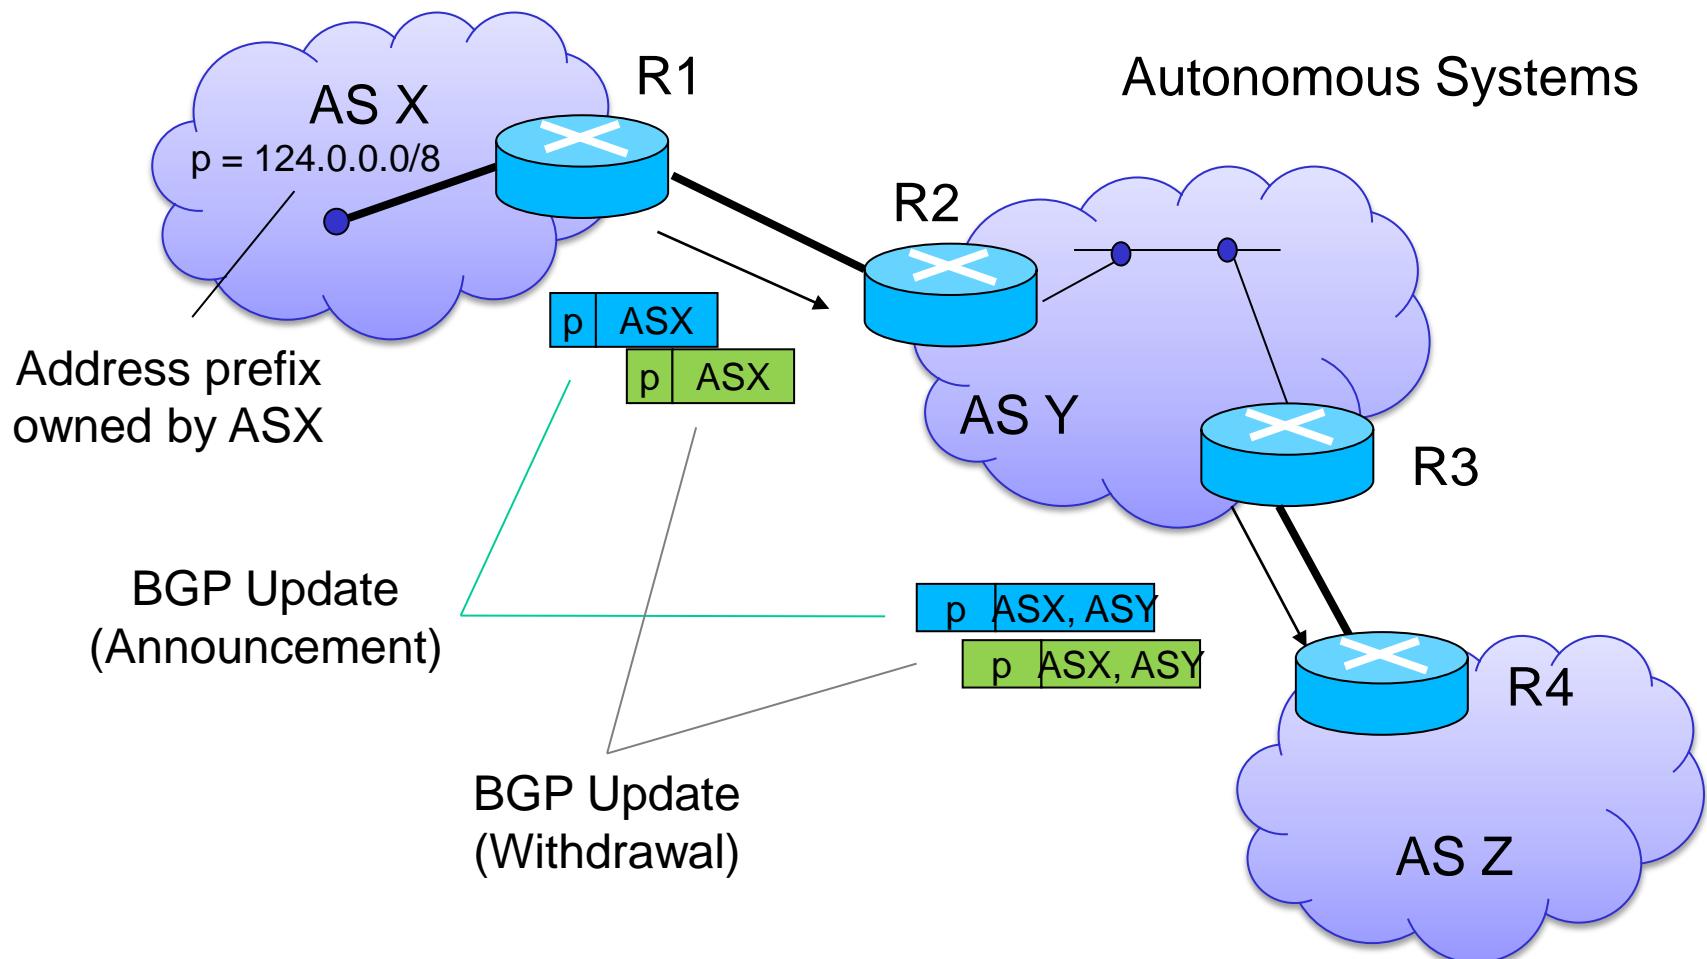
AS-CRED

- Behavior Analysis
- Reputation Computation
- Alert Generation

Performance Analysis

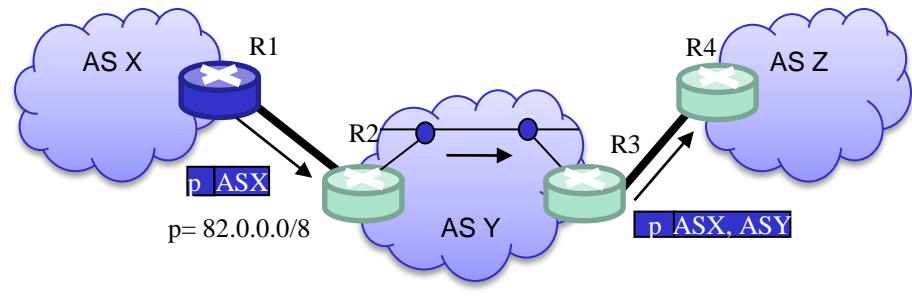
Conclusion and Future Work

Border Gateway Protocol



Problems: Inaccurate BGP Updates

- Announcement of IP prefixes *not owned* by ASX or are *bogons*
- Persistent and well-known problem
- Reasons for occurrence:
 - Blocking Content**
 - YouTube was unavailable for about 1 hour when its Prefix was hijacked by Pakistan Telecom AS 17557
 - Spamming**
 - AS 8717, an ISP in Sofia, Bulgaria, originated announcements for 82.0.0.0/8
 - May due to malicious intent or misconfiguration



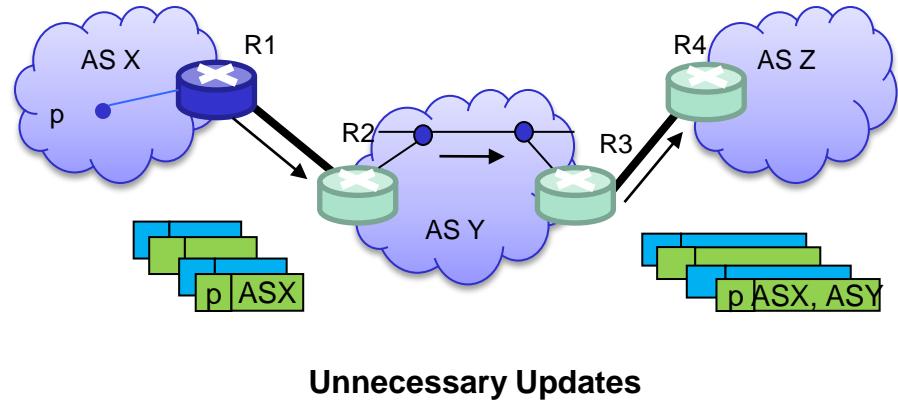
Inaccurate Updates

Well-known Incidences

Prefix hijacked	Victim AS	Attacker AS	Dates
63.218.188.0/22	3491	23724	April 8, 2010
194.9.82.0/24	36915	6461	March 15, 2008
208.65.153.0/24	36561 (YouTube)	17557	Feb. 24, 2008
66.135.192.0/19	11643 (ebay)	10139	November 30, 2007
12.0.0.0/8	7018	31604	Jan. 13, 2007
82.0.0.0/8	NULL	8717	Dec. 2004 - Jan. 2005
61.0.0.0/8	4678	17607	Dec. 2004 - Jan. 2005

Problems: Unnecessary BGP Updates

- Repeated announcement and withdrawal of IP prefixes *owned* by ASX, or illegal AS values in update message
- Persistent and *NOT* well-known problem
- Order of magnitude larger problem compared with prefix hijacking
- Principal suspected reason – Misconfiguration of BGP router
- Example:
 - Prefix 41.222.179.0/24 announced and withdrawn 4824 times by AS37035 between Dec. 3, 2009 and Dec. 7, 2009, once every 1.5 minutes.
 - Announcement of private AS numbers (e.g., AS65535) due to improper export policy – filtering

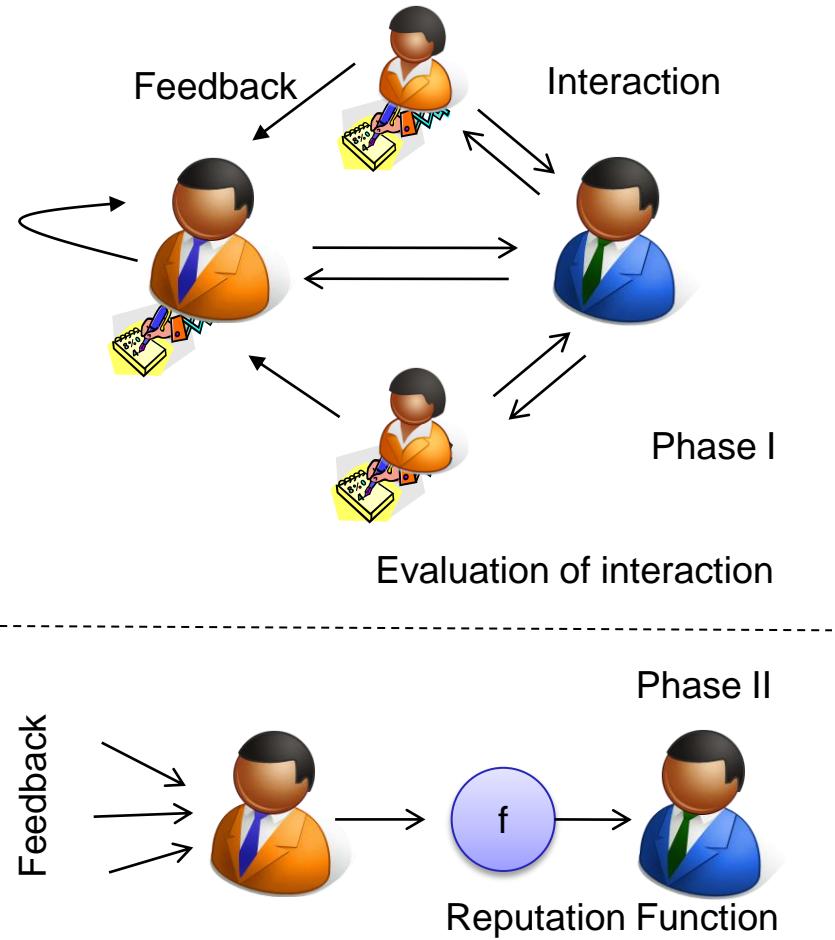


Prominent Incidences

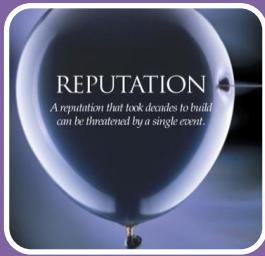
AS	Prefix	Dates	RAW
7035	41.222.179.0/24	Dec. 3 – Dec. 7 2009	4824
8452	41.235.83.0/24	Nov. 2 - Nov. 10, 2009	2088
704	152.63.49.180/30	Dec. 8 - Dec. 31, 2009	1628
145	140.217.157.0/24	Nov. 1 - Nov. 27, 2009	1080

Approach

- Principal Question:
 - How do we know if ASes are announcing valid updates ?
 - Update Validity: necessary and accurate
- Approach:
 - Essentially a question of **trust** – a subjective expectation on the behavior of an entity
 - In this problem:
 - Entity – Autonomous Systems
 - Behavior – announcement of valid BGP updates
- Observation:
 - ASes repeat their behaviors
 - Past can be used to predict future
 - Metric of choice: **Reputation**



Goals



Compute the reputation for Autonomous Systems in the Internet, by analyzing past BGP updates announced by them for their validity – accuracy and necessity.

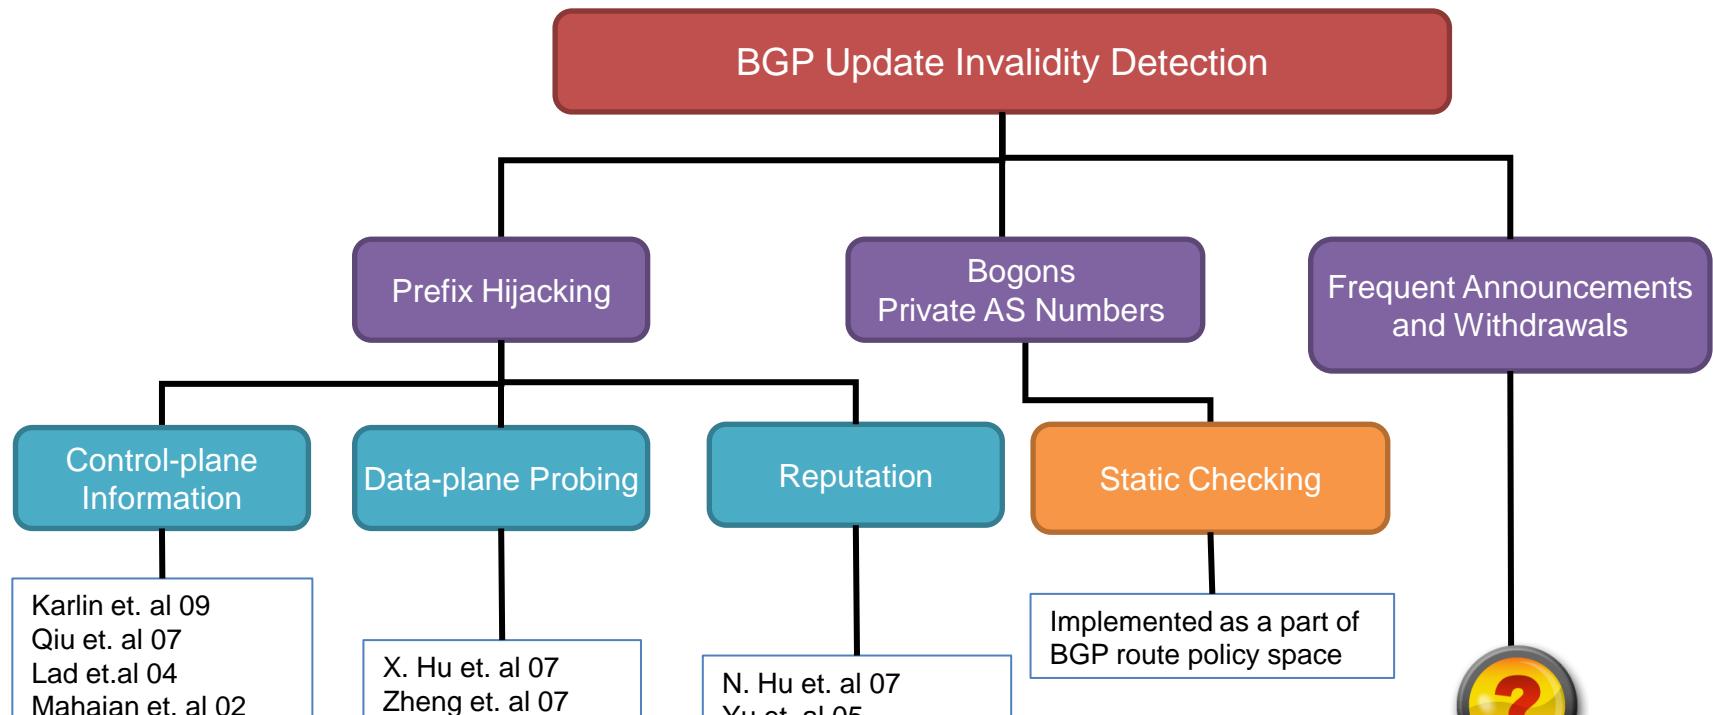


Provide an alert service for tracking the subsequent announcement of potentially invalid BGP updates based on the computed reputation.



Deploy as an publicly available service for everyone to use.

Traditional Approach

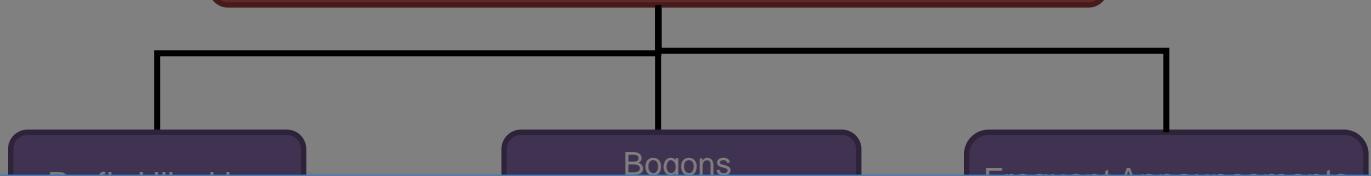


- Use Short-lived prefix announcements as basis for detection
- Consider them both malicious and misconfigured
- Provide alerts for potential hijacks

- Third-Party Feedback Dependent
- Requires Overlay Trust Network

Traditional Approach

BGP Update Invalidity Detection



Principal Issues:

- No Non-necessity check
- No quantitative modeling of AS behavior tendencies
 - High False Positives

Lad et.al 04

Mahajan et. al 02

Xao et. al 02

X. Hu et. al 07

Zheng et. al 07

Zhang et. al 05

N. Hu et. al 07

Yu et. al 05

BGP Route policy space

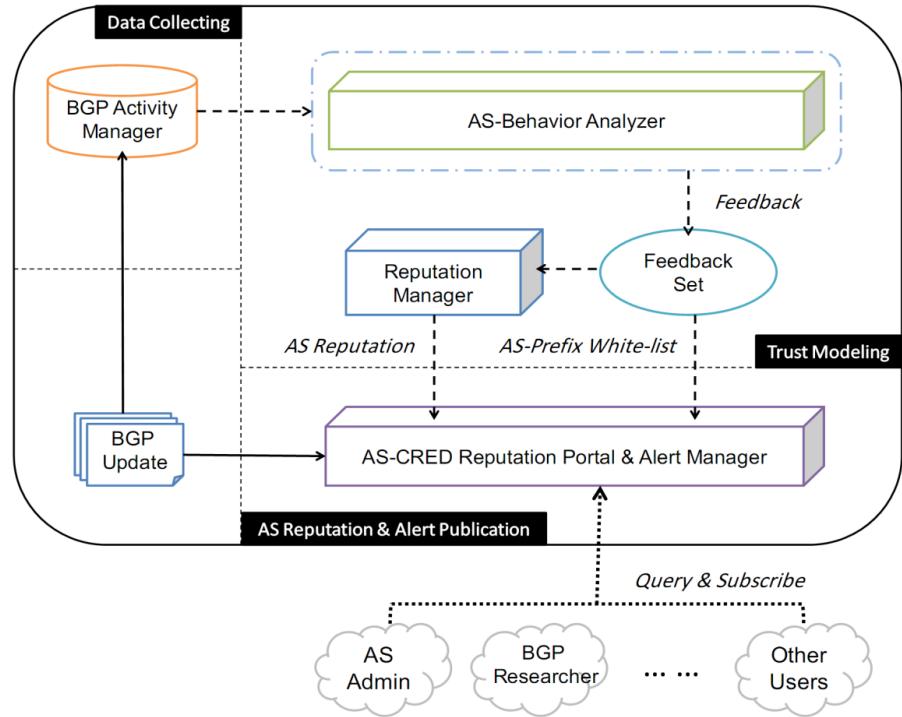
- Use Short-lived prefix announcements as basis for detection
- Consider them both malicious and misconfigured
- Provide alerts for potential hijacks

- Third-Party Feedback Dependent
- Requires Overlay Trust Network



AS-CRED: Architecture

- **BGP Activity Manager:**
 - Database for BGP updates
 - Obtained from well-connected BGP data collectors
- **AS-Behavior Analyzer:**
 - Analyzes the updates in BGP Activity Manager, based on a set of well-defined properties to **detect** invalidity
 - The results of the analysis, is a feedback on the past behavior of ASes
- **Reputation Manager:**
 - Computes the reputation of the ASes based on a well defined mathematical function
 - Uses past behavior information in the form of feedback
- **Reputation Portal:**
 - Once the AS reputations are computed it is made available through a web portal
- **Alert Manager:**
 - Uses AS reputation, to trigger real-time alerts regarding potential invalidity of any new updates propagated within the Internet.



AS-CRED Architecture

Data Source: RouteViews

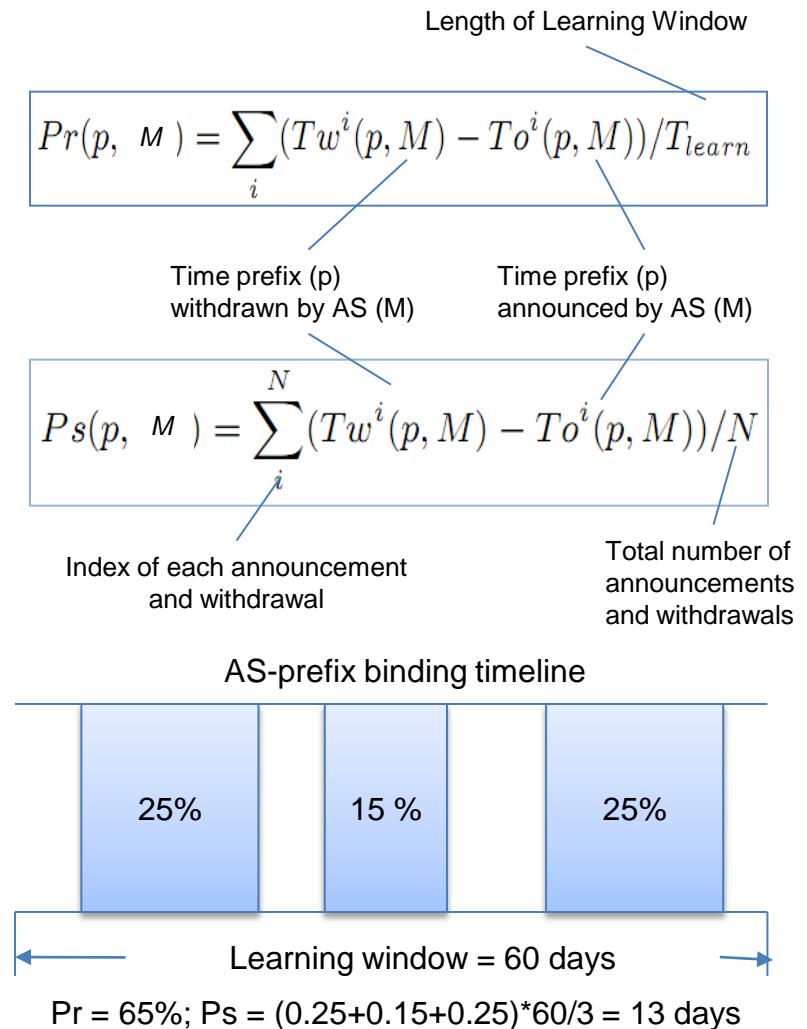
- Basically a group of BGP routers (AS 6447) peered with about 40 other ASes at crucial places
- Receives updates from the peers which it stores in its database without any filtering
- Maintains RIB dumping database: a **prefix list** with time-stamped information on **origin** and **AS-path**
- Route-Views does not originate any prefix or forward a received update message
- RIB dumping **every two hours**, update messages every **15 minutes**
- Useful for analyzing past behaviors of ASes



For every prefix visible to ASes X, Y and Z an entry exists in 6447

Behavior Analysis: Property I

- **Observation:** AS-prefix bindings which are invalid usually last for a short period of time, *i.e.*, they are unstable.
- **Aim:** Detect AS-prefix bindings stability
- **Need:** Historical Information based analysis
 - Analysis window (60 days learning window)
 - Two complimentary metrics
 - **Prevalence** – percentage of learning window AS-prefix binding lasted
 - **Persistence** – average time an AS-prefix binding lasted



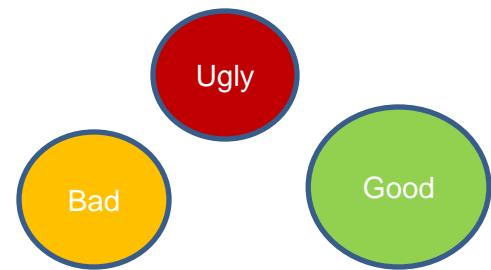
Property II & Feedback

Initial Classification

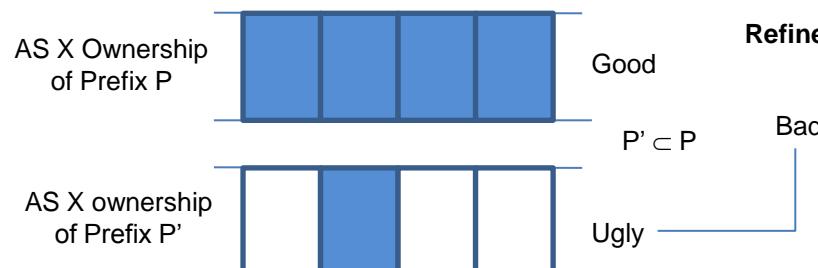
Prevalence	Persistence	Feedback
Hi	Hi	Good
Hi	Lo	Bad (Unnecessary)
Lo	Hi	Good
Lo	Lo	Ugly (Inaccurate)

Entry format

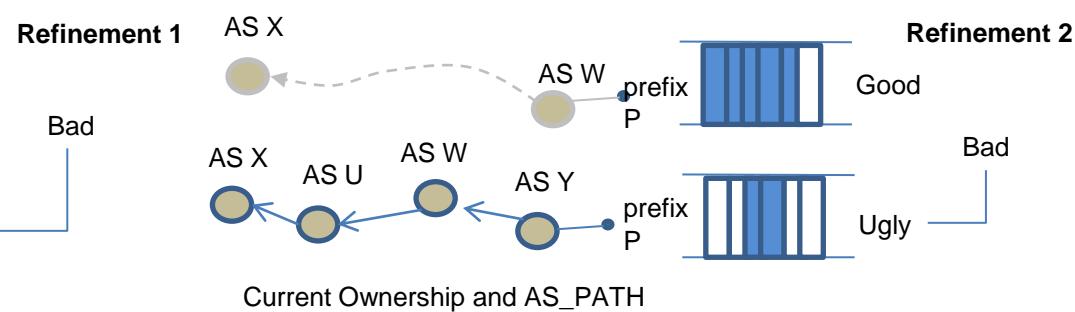
AS	prefix	Timestamp of announcement
----	--------	---------------------------



Refinement

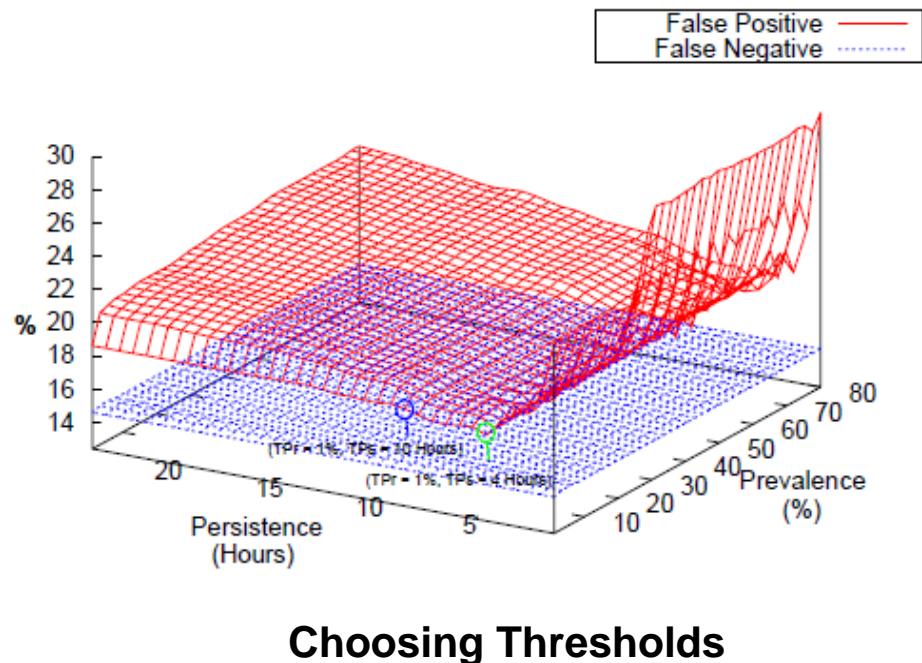


Refinement 1



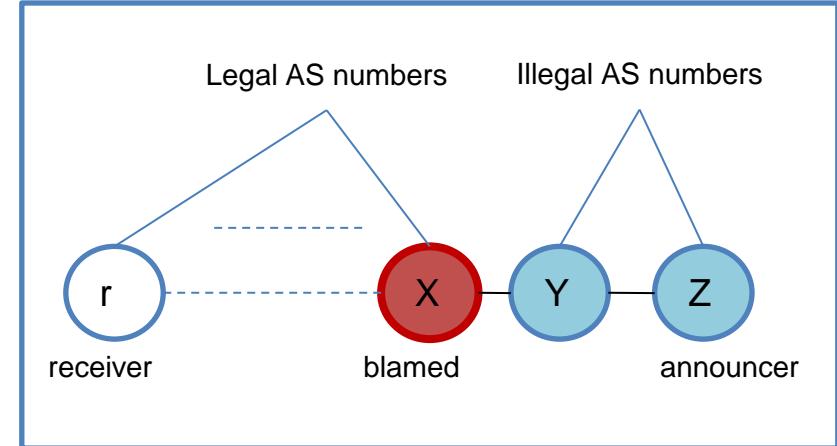
Stability Threshold

- Feedback results in three sets:
 - Good, Bad and Ugly
- Threshold needed to determine:
 - What is Hi and Lo ?
- Generated based on comparison with Internet Route Registries (IRR), the closest source to ground truth available
- Compare
 - **False Positive**: entries in IRR found in Ugly set
 - **False Negative**: entries not in IRR found in Good and Bad set
- Value of choice: $T_{Pr} = 1\%$ and $T_{Ps} = 10$ hours

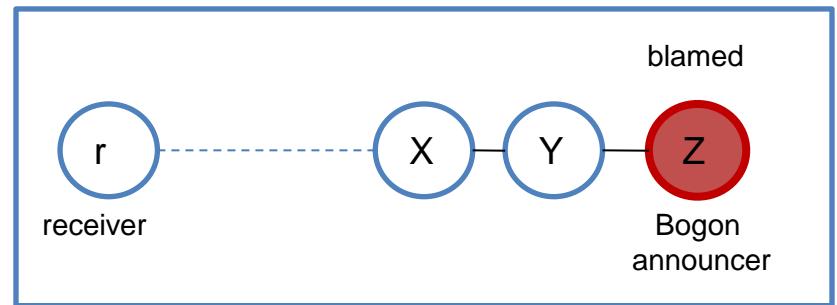


Behavior Analysis: Property II

- **Observation:** BGP updates contain illegal values for ASes and the prefixes they announce
 - **Illegal AS numbers:**
 - Example, those in the range of: 64496-64511, 64512-65534
 - **Bogons:**
 - Set of yet to be allocated prefixes
- **Feedback:**
 - **Illegal AS numbers:**
 - First AS in the AS-PATH with a legitimate value blamed
 - Update considered **Unnecessary**
 - **Bogons:**
 - The announcer is blamed
 - Update considered **Inaccurate**



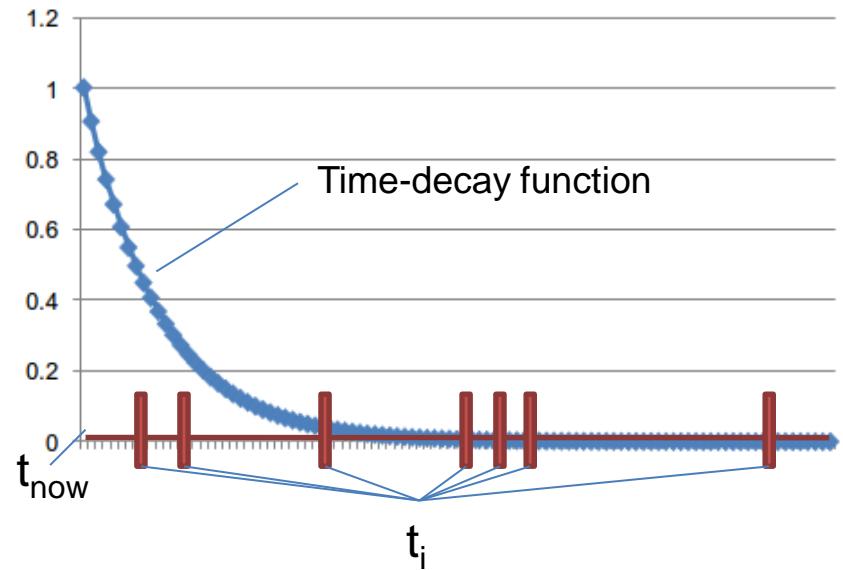
Illegal AS Number



Bogon Announcement

Reputation Computation

- AS-CRED computes
 - **untrustworthiness** of ASes in announcing valid updates
 - Reputation of an AS is computed based on Bad and Ugly feedback only
- Uses a time-decay function where
$$Rep_X(a) = \sum_{t_i} 2^{-(t_{now} - t_i)/h_X}$$
 - X is either B or U
 - h_X is a half-life of behavior X
 - t_{now} is the current time
 - t_i is the feedback timestamp:
- Two reputation values created for each AS
 - **RepU** – characterizes an AS's past inaccurate update announcement
 - **RepB** - characterizes an AS's past unnecessary update announcement

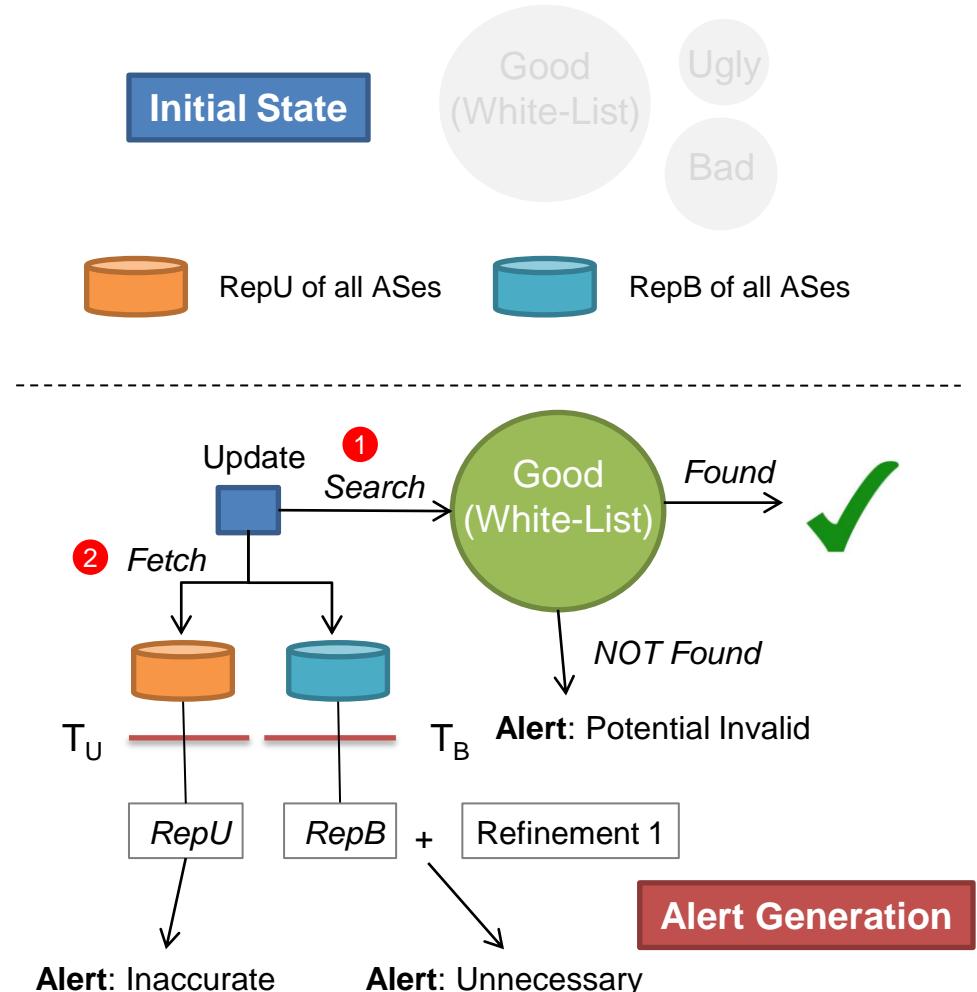


- **Half-life:** time by which the weight of the reputation of an AS is halved
- Set based on when 75% of the ASes repeat their invalid updates
- Values: $h_U = 3 \text{ days}$, $h_B = 6 \text{ days}$

Alert Generation Process

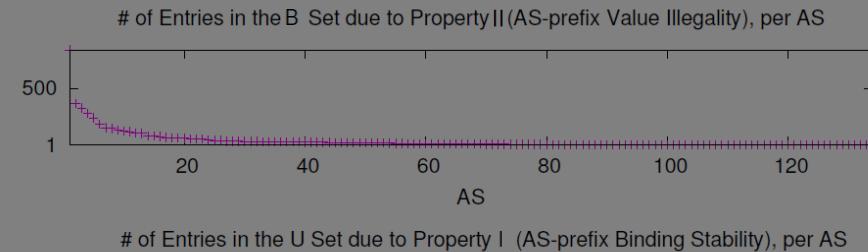
Three Steps Process

- White-List Filtering:**
 - When a new update is received, we first checks to see if its corresponding AS-prefix binding (a, p) is in our white-list (G set)
- Alert Generation:**
 - If (a, p) are not in the white-list, we post an *potential invalid* Alert
- Relabeling:**
 - Label updated to *Unnecessary*, if
 - $RepB(a)$ is poor or $RepU(a)$ is poor with $p \subset p'$ such that (a, p') is in the white-list.
 - Label updated to *Inaccurate*, if
 - $RepU(a)$ is poor with no $p \subset p'$ such that (a, p') is in the white-list



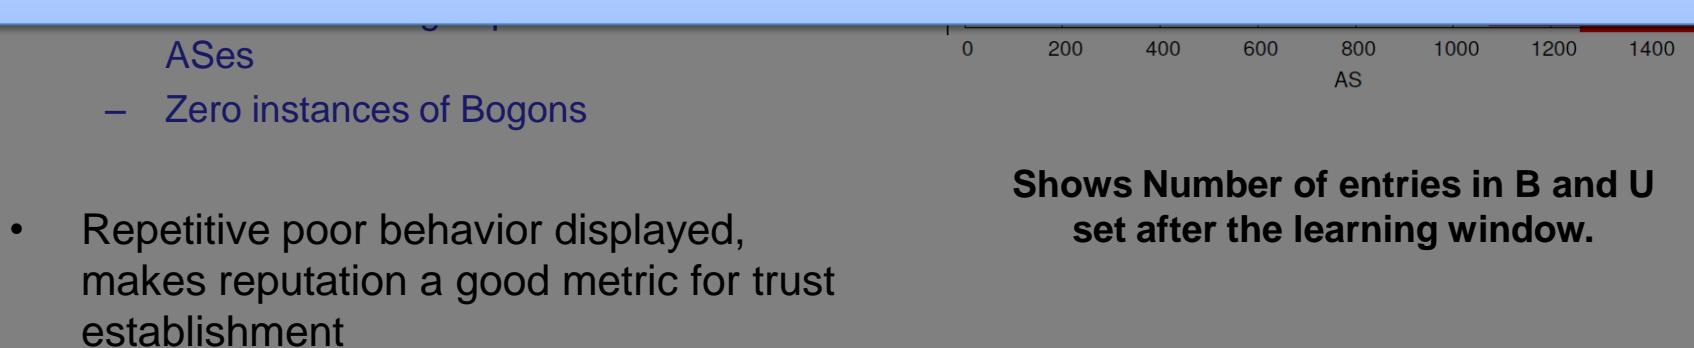
Behavior Analysis (Nov 1, '09- Dec 30, '09)

- Property I:
 - Unnecessary repeated updates far outnumber prefix hijackings or updates with illegal AS numbers
 - Updates for prefix hijacking and illegal AS numbers instances are similar in scale



Observation:

- Unnecessary updates a bigger problem in inter-domain routing compared to updates with Inaccurate information



Quality of Behavior Analysis

- Inaccurate Updates
 - U set stores instances of inaccurate updates –prefix hijacking
 - Inaccurate updates detected compared with *Internet Alert Registry* w.r.t. IRR
 - **4 fold improvement in False Positives**
- Unnecessary Updates
 - B set stores instances of Unnecessary updates
 - Unnecessary updates from repeated announcements and withdrawals were
 - 92% legitimate AS-prefix bindings (based on Internet Route Registry)
 - Announced 42 times more often than Good AS-prefix bindings

Scheme	No Record	False Positive		Hijack
		IRR Match	No IRR Match	
AS-CRED	841 (13.7%)	975 (18.4%)	4323 (81.6%)	
IAR	4190 (10.7%)	25892 (74.4%)	8903 (25.6%)	

Behavior Analysis (Nov 1- Dec 30) Vs. IAR w.r.t. IRR

# Announcements and Withdrawals			
AS	Prefix	NAW	Duration Observed
8452	41.235.83.0/24	2088	Nov 2- 10, 2009
704	152.63.49.180/30	1628	Dec 8 – 31, 2009
145	140.217.157.0/24	1080	Nov 1- 27, 2009

Prominent Examples of Unnecessary Updates

Behavior Analysis Overall Statistics

Prefix Statistics

Property	Value
Prefixes Observed	367605
SOAS Prefix Observed	357855
MOAS Prefix Observed	9750

AS Statistics

Property	Value
AS Observed	33925
AS announcing Unnecessary Updates	1568 (4.6%)
AS announcing Inaccurate Updates	693 (2.0%)
AS exclusively announcing Unnecessary Updates	79
AS exclusively announcing Inaccurate Updates	89

AS-Prefix Binding Classification

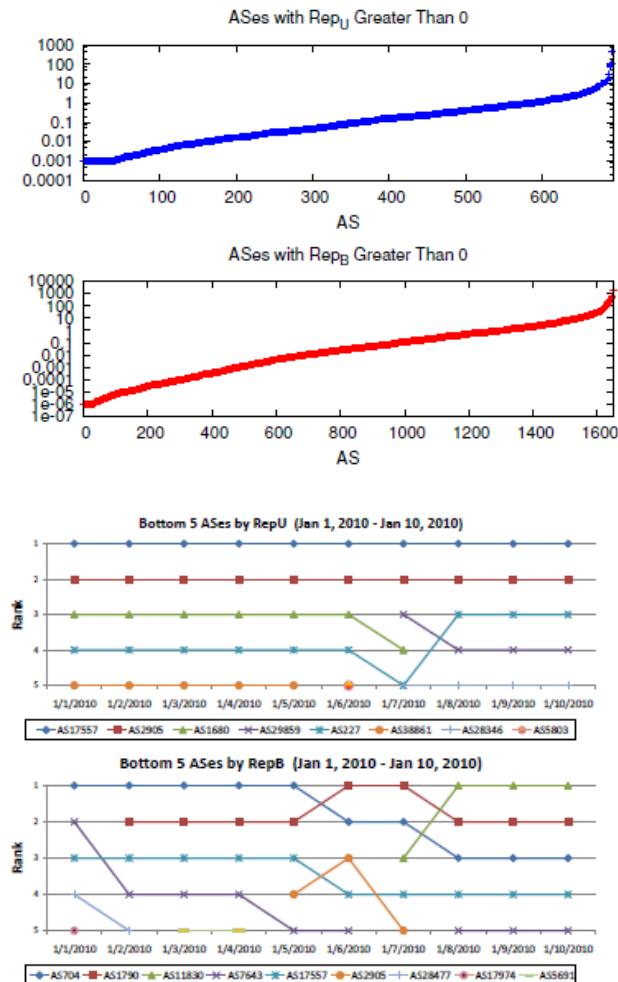
Property	Value
Total AS-Prefix Bindings	376224
AS-Prefix Bindings in Inaccurate Updates	6139
AS-Prefix Bindings in Unnecessary Updates	26270

Behavior Incidences Statistics

Property	Value
Number of Inaccurate Updates	13615
Number of Unnecessary Updates	213725

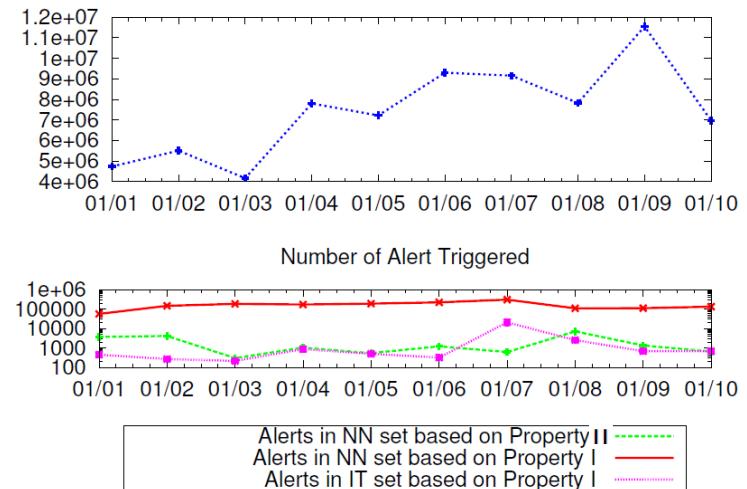
Reputation Analysis

- AS-CRED Reputation characterizes the current perpetrators of invalid updates announcement:
 - ZERO reputation is considered good behavior
 - 693 ASes have $Rep_U > 0$
 - 1568 ASes have $Rep_B > 0$
 - 90% of ASes with poor behavior have reputation close to ZERO
- ASes show repetitive behaviors
 - Most ASes are good, very few ASes demonstrate repeated poor behaviors
- AS-CRED is **sensitive** in detecting even announcers of one-off invalid updates



Alert Consistency

- Given AS reputation, newly received updates received over Jan 1, 2010 – Jan 10, 2010 are be evaluated
- Updates not seen in white-list classified as unnecessary or inaccurate based on reputation of announcing AS
- Sets
 - IT - stores all inaccurate updates
 - NN - stores all unnecessary updates
- We use 60 day consistency check window (Nov 20, 2009-Jan 20, 2010) to:
 - Determine if the prediction was accurate
 - Based on behavior analysis



Classification	Count
Total NN set entries	3546
NN set entries classified in G set	71 (2.5%)
NN set entries classified in B set	2591 (97.4%)
NN set entries classified in U set	3 (0.1%)
Total IT set entries	625
IT set entries classified in G set	7 (0.2%)
IT set entries classified in B set	0 (0%)
IT set entries classified in U set	618 (98.8%)

Alert Accuracy

- For updates deemed inaccurate:
 - AS-CRED detects prefix hijacking in two places:
 - Behavior analysis to populate U set
 - Alert generation when RepU is used to determine if update is a hijack
 - Behavior Analysis shown to be accurate
 - Compared the alert results with Internet Alert Registry and IRR (comparative ground-truth)
 - 8 fold improvement in False Positives

Scheme	No Record	IRR Match	No IRR Match
AS-CRED	112 (18.1%)	42(8.3%)	465 (91.7%)
IAR	413 (11.2%)	2437(75.4%)	798 (24.6%)

Alert Generation (Jan 1-Jan 10) vs. IAR w.r.t. IRR

- For updates deemed unnecessary :
 - 88% of the associated AS-prefix binding found in IRR
 - Average NAW – 26 with the maximum 4492
 - Contrast for AS-prefix binding in Good set (Avg. NAW ~ 1)

AS-CRED Service Screenshot

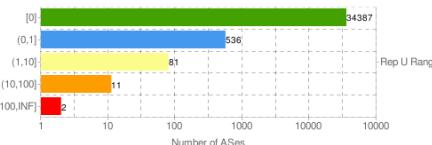
WORST PERFORMERS AND REPUTATION STATISTICS

DATE: 19-MAY-2010

ASes with Worst RepU

AS Number	Reputation U
23724	292.55
17557	207.45
10318	48.31
6298	26.27
2072	22.06

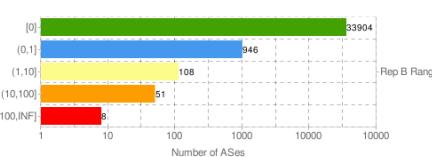
RepU Histogram



ASes with Worst RepB

AS Number	Reputation B
4538	3336.73
5769	703.39
668	257.84
17557	255.29
8551	235.75

RepB Histogram



AS REPUTATION SEARCH

SEARCH AS 17557

RepU of AS 17557

207.45



90 Days RepU Trend



RepB of AS 17557

255.29



90 Days RepB Trend



Bottom 5 ASes by Reputation

Past Reputation Trend for an AS

Reputation-based Update Alert

<http://rtg.cis.upenn.edu/qtm/ascred/>

POTENTIALLY INVALID BGP UPDATES

(Latest 5 listed. List updated every hour.)

ASN	IP Prefix	Time	Alert Type	RepU	RepB
491	134.135.16.0/20	Wed May 19 23:56:03 2010	Pot. Invalid	0.16	5.89
491	209.22.19.0/24	Wed May 19 23:56:03 2010	Pot. Invalid	0.16	5.89
491	140.175.0.0/16	Wed May 19 23:56:03 2010	Pot. Invalid	0.16	5.89
491	209.22.18.0/24	Wed May 19 23:56:03 2010	Pot. Invalid	0.16	5.89
491	209.22.20.0/24	Wed May 19 23:56:03 2010	Pot. Invalid	0.16	5.89

LIKELY INVALID BGP UPDATES

(Latest 10 listed. List updated every hour.)

ASN	IP Prefix	Time	Alert Type	RepU	RepB
17557	119.73.35.92/32	Wed May 19 23:33:13 2010	Inaccurate	207.45	255.29
17557	116.71.180.106/32	Wed May 19 23:32:11 2010	Unnecessary	207.45	255.29
17557	119.152.21.241/32	Wed May 19 23:01:59 2010	Unnecessary	207.45	255.29
17557	119.73.35.38/32	Wed May 19 22:03:05 2010	Inaccurate	207.45	255.29
17557	124.29.192.54/32	Wed May 19 19:34:05 2010	Inaccurate	207.45	255.29

Conclusions & Future Work

- **Conclusions:**

- *Repetitive Behavior*: ASes which announce invalid updates do so repeatedly, which makes reputation a good metric to characterize them
- *Large number of Unnecessary Updates*: The number of unnecessary updates with poor stability far outnumber the inaccurate ones and those with illegal values
- *Sensitivity*: The reputation metric is very sensitive and can capture ASes which seldom announce invalid updates
- *Improved Hijack Detection*: The AS-behavior analysis and alert service are much more accurate than existing services (such as the IAR) for detecting prex hijacking
- *Consistency of Analysis and Reputation*: The reputation assigned to an AS is a representative and behavior predictive value.

- **Future Work:**

- Extend this work by including other properties for determining an AS' tendency to announce valid updates, such as presence of valley-free path and stable links in the AS-PATH.

Thank You & Questions

